

COMPUTING SUBFIELDS OF NUMBER FIELDS AND APPLICATIONS TO GALOIS GROUP COMPUTATIONS

ANDREAS-STEPHAN ELSENHANS AND JÜRGEN KLÜNERS

ABSTRACT. A polynomial time algorithm to give a complete description of all subfields of a given number field was given in [HKN].

This article reports on a massive speedup of this algorithm. This is primary achieved by our new concept of *Galois-generating subfields*. In general this is a very small set of subfields that determine all other subfields in a group-theoretic way. We compute them by targeted calls to the method from [HKN]. For an early termination of these calls, we give a list of criteria that imply that further calls will not result in additional subfields.

Finally, we explain how we use subfields to get a good starting group for the computation of Galois groups.

1. INTRODUCTION

Given a number field L/\mathbb{Q} of degree n , we can ask for a description of all fields K such that $\mathbb{Q} \subset K \subset L$. It turns out, that the number of these fields is not bounded by a polynomial function in the degree of L/\mathbb{Q} . Thus, there can not be a polynomial time algorithm for this task.

However, there are less than n subfields K_1, \dots, K_m , such that all other subfields are intersections of some of these fields. Any list of subfields with this property is called a list of intersection-generating subfields. Initially, these fields were called generating subfields, but we prefer to use the more precise term intersection-generating as there are other operations (e.g. composition) that can construct a subfield out of others.

A subfield $\mathbb{Q}[\alpha]$ of L is described by the minimal polynomial of α and an embedding $\alpha \mapsto h(\beta) \in L$ for some polynomial $h \in \mathbb{Q}[Y]$. Here, β denotes the primitive element of L that is used to represent L . The fundamental new idea of [HKN] is that we can find the subfield as a vector space by using the LLL-algorithm [LLL]. From that we derive a primitive element α , determines its minimal polynomial and the embedding. At most $(n - 1)$ of these steps are necessary to compute all so called principal subfields. Note, that a list of all principal subfields is intersection-generating.

The runtime depends heavily on the number of LLL-calls. In this note we will show how hard cases ($n = 60, \dots, 100$) can be done with very few (usually < 10) LLL-calls. The subfields found will not be intersection-generating. But, they will be *Galois-generating*. This is the new term that we introduce in this article.

An other element of our improvements is that we skip LLL-calls that would reproduce subfields that are already known. To detect this, we give a list of criteria in Algorithm 8.1.

All algorithms are of a p -adic nature. That means we work with p -adic root approximations in a p -adic splitting field of f . The prime p is chosen by the algorithms.

To derive the other subfields from the Galois-generating ones we have to compute the intersection of certain wreath products. For this we give a graph-theoretic algorithm in section 5. The resulting group will be an overgroup of the Galois group of L . Thus, it can be used as a starting group for the computation of the Galois group of L by the Stauduhar [St] method. In section 9 we give an algorithm to refine this starting group.

All the algorithms described here will be available in MAGMA 2.23 [BCP].

Composition-generating subfields. One could also ask for composition-generating subfields. This seems to be a challenging question that is almost independent of this investigation.

Acknowledgement. Many thanks to N. Sutherland for writing the MAGMA implementation of the method from [HKN] that served as a starting point of this investigation.

2. NOTATION

In this article, we will consider the number field L as an extension of the rationals. The letter n denotes the degree of L . The minimal polynomial of a chosen primitive element β of L will be denoted by f . For simplicity we assume the roots of f to be algebraic integers. We denote approximations of the roots by r_i . Further, we denote the normal closure of L by \tilde{L} .

The letter K denotes a subfield of L . The minimal polynomial of a primitive element α is called g . In the case that we are dealing with several subfields we add indices to K , α , and g .

We will use the term Galois group in the following way: The Galois group of a polynomial f is the Galois group of a splitting field represented as a permutation group acting on the roots of f . The Galois group of a number field is the Galois group of the minimal polynomial of a primitive element.

Any permutation group that is an overgroup of the Galois group of f is called a *Galois starting group* of f .

When we give examples that involve permutation groups, we write nTk for the k -th transitive group of degree n in the database of transitive groups [Hul, CH].

3. GALOIS THEORY OF SUBFIELDS AND BLOCK SYSTEMS

Recall 3.1. Let G be a transitive permutation group of degree n . We will use the following terms and facts [DM, HEB, Hup].

- (1) A system of blocks (or partition) $\mathcal{B} = \{B_1, \dots, B_k\}$ is a set of disjoint, non-empty subsets of $\{1, \dots, n\}$ that cover all of $\{1, \dots, n\}$ such that $\sigma B_i \in \mathcal{B}$ for all $\sigma \in G$ and $i = 1, \dots, k$.
- (2) Let \mathcal{B}_1 and \mathcal{B}_2 be two systems of blocks. In the case that each block of \mathcal{B}_1 is a subset of a block of \mathcal{B}_2 , we call \mathcal{B}_1 finer than \mathcal{B}_2 .
- (3) A system of blocks is the same as the equivalence classes of a G -equivariant equivalence relation. The relation simply encodes which pairs (r_i, r_j) are in the same block.

- (4) A system of blocks \mathcal{B} is called *principal*, if there is a pair (i, j) in one block of \mathcal{B} , such that \mathcal{B} is the finest system of blocks, such that i and j are in one block.

As G is transitive, we can set $i = 1$ without loss of generality. Thus, there are at most $(n - 1)$ principal partitions.

- (5) The equivalence relation corresponding to a principal partition containing (i, j) is the reflexive, transitive and symmetric closure of $\{\sigma(i, j) \mid \sigma \in G\} \subset \{1, \dots, n\}^2$.
- (6) We say that the pair (i, j) generates the relation and the corresponding principal system of blocks.
- (7) Let \mathcal{B} be a system of blocks and $B_1 \in \mathcal{B}$ be the block containing 1. Then the principal partitions that are finer than \mathcal{B} are exactly those that are generated by $(1, i)$ for $i \in B_1, i \neq 1$.
- (8) Let $\mathcal{B}_1, \dots, \mathcal{B}_r$ be some systems of blocks. The finest system of blocks \mathcal{B} such that all the \mathcal{B}_i are refinements of \mathcal{B} corresponds to the equivalence relation that is the transitive closure of the union of the relations corresponding to the \mathcal{B}_i for $i = 1, \dots, r$.
- (9) There is a 1–1 correspondence between the super-groups of the stabilizer of 1 and the block systems. It is given by mapping a system of blocks to the stabilizer of the block containing 1.
- (10) The largest subgroup of $S_{\ell k}$ that has a system of ℓ blocks of size k is the wreath product $S_k \wr S_\ell$. It is isomorphic to the semi-direct product $S_k^\ell \rtimes S_\ell$. The action on S_k^ℓ is given by permutation of the components.

Remark 3.2. The main theorem of Galois theory gives us a 1–1 correspondence between the intermediate fields of a normal closure of a field extension and the subgroups U of its Galois group. Two subgroups $U_1 \subset U_2$ correspond to two subfields $K_2 \subset K_1$. I.e., the correspondence is inclusion-reversing.

We view the Galois group of L as a subgroup of S_n that acts on the roots of f . Then the stem field L corresponds to the stabilizer of the first root. This shows that we have a 1–1 correspondence between the block systems of G and the intermediate fields of L/\mathbb{Q} .

We can use the following algorithm to compute all $\text{Gal}(\tilde{L}/\mathbb{Q})$ -equivariant relations on the set of roots of f that correspond to principal partitions.

Algorithm 3.3 (Principal systems of blocks by 2-set resolvent).

- (1) Find a resolvent Res_2 representing the 2-sets of roots of f . We can use $\prod_{i < j} (X - (r_i + r_j))$ or $\prod_{i < j} (X - r_i r_j)$ in the case that one of them does not have multiple zeros. In general, we have to find a shift $s \in \mathbb{Z}$, such that $\prod_{i < j} (X - (r_i + s)(r_j + s))$ is squarefree.
- (2) Factor the resolvent Res_2 .
- (3) For each irreducible factor F of Res_2 , identify which pairs of roots are encoded in it. I.e., determine $R_F := \{(i, j) \mid F((r_i + s)(r_j + s)) = 0\}$.
- (4) For each relation R_F , compute the reflexive, transitive, and symmetric closure.
- (5) Return a list of all equivalence relations found.

Remarks 3.4. (1) Assume the above algorithm results in the G -equivariant equivalence relations R_1, \dots, R_k . Then all other G -equivariant equivalence

relations are given as the transitive closure of $R_{i_1} \cup \dots \cup R_{i_m}$, for any choice of indices $i_1, \dots, i_m \in \{1, \dots, k\}$.

- (2) From this we see, that the subfields corresponding to the R_i are intersection-generating subfields. As G acts transitively on G , there is no orbit on 2-sets shorter than $\frac{n}{2}$. Thus, we reproduce the bound of $(n-1)$ intersection-generating subfields.
- (3) It is known that the computation of a 2-set resolvent can be done in polynomial time [BFSS] and polynomial factorization is polynomial time as well [LLL]. Thus, we have a polynomial time algorithm to compute intersection-generating subfields as soon as we can find a p -adic splitting field of polynomial complexity in polynomial time.

A skillful way to circumvent too large splitting fields is given in [Kl, p. 255]. The basic idea is to work in the smallest field which is sufficient to represent the polynomial $\prod_{i \in B} (X - r_i)$ for a block B . Note that, the degree of the field of definition of this polynomial is bounded by the degree of the subfield we are looking for.

Thus, we have a polynomial time algorithm that computes all principal subfields.

- (4) In case we ask for the list of all subfields instead of a list of intersection-generating subfields we can either compute the intersections directly, or we can use the description of the subfields in terms of systems of blocks. This means, we have to list all equivalence relations that result as the transitive closure of the union of some equivalence relations already known.

Assume the roots r_1, \dots, r_n of f are given in a (p -adic) splitting field of f , we need algorithms to make the correspondence between subfields and block systems explicit. For this we first introduce block invariants.

Definition 3.5. Let M be a nontrivial subset of $\{1, \dots, n\}$. A polynomial $I \in \mathbb{Z}[X_1, \dots, X_n]$ is called a block invariant for M , if I only involves the variables $\{X_k : k \in M\}$ and is invariant under all permutations of the elements of M .

Example 3.6. A nontrivial block invariant of smallest degree is given by $\sum_{i \in M} X_i$. More generally, we can use $\sum_{i \in M} T(X_i)$ with any transformation $T \in \mathbb{Z}[X]$.

Other block invariants are given by $\prod_{i \in M} X_i$ and $\prod_{i \in M} (X_i + s)$ for any $s \in \mathbb{Z}$.

Remark 3.7. In the language of invariant theory, I is also called a relative invariant for $\text{Stab}_{S_n}(M) \subset S_n$. The interested reader may consult [E2] for more constructions of relative invariants.

Remark 3.8 (Block system to a subfield). Let $L/K/\mathbb{Q}$ be a tower of fields. Further, let $h(\beta)$ be a primitive element of the degree m subfield K .

In this situation r_i and r_j are in the same block (of the block system corresponding to K) if and only if $h(r_i) = h(r_j)$. In the case that we work with root approximations we can not prove $h(r_i) = h(r_j)$. But, we are able to disprove the equality. As the total number of blocks is m and each block has n/m elements, this suffices to determine the block system.

The explicit description of the inverse operation is more complicated as we have to compute the minimal polynomial of a primitive element α of the subfield and its image in \tilde{K} .

Algorithm 3.9 (Subfield from block system). *Given a system of blocks $\mathcal{B} = \{B_1, \dots, B_m\}$ for the p -adic root approximations r_1, \dots, r_n of f in some extension of \mathbb{Q}_p . We assume the r_i to be algebraic integers. This algorithm computes a defining polynomial g for the corresponding subfield.*

- (1) Find a non-degenerated block invariant $I = \prod_{i \in B_1} (r_i - s)$ for some $s \in \mathbb{Z}$. Here, non-degenerated means, that the values $I_k := \prod_{i \in B_k} (r_i - s)$ on all blocks are pairwise distinct. In the case that $p > n^2$, we can even find an s such that these products are different in the residue class field of the splitting field [K1, Lemma 42].
- (2) Use a bound on the complex roots of f to bound the absolute values of the block invariant chosen.
- (3) Derive a bound for the coefficients of $g := \prod_{k=1}^m (X - I_k)$ and a p -adic precision that suffices to reconstruct the coefficients of g .
- (4) Compute the roots r_i , the values of I_k and the coefficients of g as p -adic numbers with the precision found.
- (5) Reconstruct $g \in \mathbb{Z}[X]$.

The embedding is given as the solution of the interpolation problem $h(r_i) = I_k$ for all $i \in B_k$ and all $k \in \{1, \dots, m\}$. Here, h is a polynomial of degree at most $(n-1)$ with rational coefficients. It defines the embedding $\iota: K \rightarrow L$ via $\iota(\alpha) = h(\beta)$.

The conditions $h(r_i) = I_k$ for all i, k with $i \in B_k$ above encode an interpolation problem that translates to a uniquely solvable linear system for the coefficients of h . To solve it efficiently, we use Newton-lifting as follows.

Algorithm 3.10 (Embedding of subfield). *Let $L/K/\mathbb{Q}$ be a tower of fields and the block invariant used to construct K be given. This algorithm computes the embedding of K in L by determining $h(\beta)$.*

- (1) Solve the interpolation problem in the residue class field of the splitting field directly. As we assume the roots to be pairwise distinct in the residue class field the system has still a unique solution.
- (2) This results in a polynomial $h_0 \in \mathbb{Z}[X]$ with $g(h_0(\beta)) \equiv 0 \pmod{p}$.
- (3) Compute the inverse of the derivative of g at this point. I.e. find $v_0 \in \mathbb{Z}[\beta]$ with $v_0 g'(h_0(\beta)) \equiv 1 \pmod{p}$.
- (4) To increase the p -adic precision of v and h we compute the sequence

$$\begin{aligned} h_n &:= (h_{n-1} - g(h_{n-1}) \cdot v_{n-1}) \pmod{p^{2^n}} \\ v_n &:= (v_{n-1} - (g'(h_{n-1}) \cdot v_{n-1} - 1)v_{n-1}) \pmod{p^{2^n}} \end{aligned}$$

- (5) In general, only $h \cdot f'(\beta)$ is known to be an element of the equation order $\mathbb{Z}[\beta]$. Thus, we compute $h_n \cdot f'(\beta) \pmod{p^{2^n}}$. In the case that all coefficients can be represented by integers of absolute value (significantly) smaller than $\frac{1}{2}p^{2^n}$ (e.g., smaller than $10^{-6}p^{2^n}$), we guess this representative coincides with $h \cdot f'(\beta)$. In this case, a single division results in a guess for h . Otherwise we continue the lifting process.
- (6) As soon as we have a guess for h we check it by testing $g(h(\beta)) = 0$. If this check is passed, we return h as the embedding. Otherwise, we continue the lifting process.

Remark 3.11. When computing the final check $g(h(\beta)) = 0$ a surprising phenomenon occurs. If the check will be passed, the guess for $h(\beta)$ is an algebraic integer.

Thus, when evaluating $g(h(\beta))$ all intermediate results will be algebraic integers as well. Thus, the denominators occurring are bounded by the index of the equation order in the maximal order. This implies that a lot of cancellations happen after each multiplication in the evaluation process.

If the guess of $h(\beta)$ is incorrect, it is most likely not an algebraic integer. Thus, we can not expect the same amount of cancellations. In practice, this check is fast in case it gets passed and slow otherwise. That explains why we do the final check only in the case that we have a good reason to believe that we are already correct.

4. USING CYCLE TYPES

Given an irreducible polynomial $f \in \mathbb{Z}[X]$ and a prime number p that does not divide the discriminant of f , we can easily compute the roots of the reduction of f modulo p in an extension of \mathbb{F}_p together with the action of the Frobenius on the roots. Is is well known, that this permutation is an element of the Galois group of f that can be identified up to conjugation. Further, by Chebotarev's density theorem each conjugacy class of the Galois group occurs for infinitely many choices of p .

If we would be able to match the roots of the reductions modulo different primes, we would be able to generate the Galois group of the splitting field heuristically.

This indicates, that we can get some information about the Galois group by inspecting the factorization of f modulo several primes. But, as we do not know how to identify roots of different reduction we have to use indirect approaches to get something out of this.

Currently we are only able to use the cycle type of the Frobenius element. This coincides with the degrees of the irreducible factors of f modulo p . Note that fixed points are always included in a cycle type.

4.1. A divisor for the order of the Galois group. Our algorithm will need a lower bound of the order of the Galois group. For this we will use the following.

Lemma 4.1. *Let (n_1, \dots, n_j) be the cycle type of an element of the transitive permutation group $G \subset S_n$. Then*

$$n \frac{\text{lcm}(n_1, \dots, n_j)}{\text{gcd}(n_1, \dots, n_j)} \mid \#G.$$

Proof. Denote by U the index n subgroup of G that stabilizes one point and by σ an element of G with cycle type (n_1, \dots, n_j) . Then σ^{n_i} has at least one fixed point. Therefore, it is contained in a subgroup conjugate to U . Thus, the order of σ^{n_i} is a divisor of $\#U$. This shows

$$\frac{\text{lcm}(n_1, \dots, n_j)}{n_i} \mid \#U.$$

Thus,

$$\text{lcm} \left\{ \frac{\text{lcm}(n_1, \dots, n_j)}{n_i} : i = 1, \dots, j \right\} = \frac{\text{lcm}(n_1, \dots, n_j)}{\text{gcd}(n_1, \dots, n_j)}$$

is a divisor of $\#U$ as well. As U is of index n in G the claim is proved. \square

The lemma above uses only one cycle type and the fact that f is irreducible to derive a divisor of the group order. In practice, we apply the lemma to several (e.g. n) primes and form the LCM of all divisors found.

To do even better we have to combine several cycles in a more subtle way. This is done by the following lemma.

Lemma 4.2. *Let $G \subset S_n$ be a p -group and $\sigma, \tau \in G$ be two elements of order p^e, p^f . Assume that the numbers of points with trivial stabilizer (i.e., with orbit of length p^e resp. p^f) in $\langle \sigma \rangle$ and in $\langle \tau \rangle$ do not coincide, then the order of G is divisible by p^{e+f} .*

Proof. Let e, f with $e \geq f$ be two integers with $e + f$ minimal such that the claim is not correct. We have $e, f > 0$ as otherwise the claim of the lemma is trivial.

First observe that $\langle \sigma \rangle$ and $\langle \tau \rangle$ are cyclic p -groups. Thus, they have a unique minimal subgroup. This show that the set of points with trivial stabilizer is unchanged when we switch to a nontrivial subgroup.

In the case that $e > f$ we take a point x with trivial stabilizer in $\langle \sigma \rangle$. This point has an orbit $M := \langle \tau \rangle x$. Now we inspect the subgroup $U := \text{Stab}_G(M)$. Obviously we have $\tau \in U$ and $\sigma \notin U$. Let k be the smallest integer such that $\sigma^{p^k} \in U$. Then we have $k > 0$ and $p^k \mid [G : U]$. In the case $k = e$ we are done. Otherwise U contains the nontrivial element σ^k of order p^{e-k} . As the points with trivial stabilizer in $\langle \sigma \rangle$ and in $\langle \sigma^{p^k} \rangle$ coincides, we can conclude $p^{f+e-k} \mid \#U$ and $p^{e+f} \mid \#G$.

In the case $e = f$ we assume that the number of points with trivial stabilizer for $\langle \sigma \rangle$ is bigger than for $\langle \tau \rangle$. Thus, there is a point x having an orbit of length p^e with respect to $\langle \sigma \rangle$ and another (shorter) orbit with respect to $\langle \tau \rangle$. We choose $M := \langle \tau \rangle x$. Now we inspect the stabilizer $U := \text{Stab}_G(M)$. We continue as above. We get $\tau \in U$ and $\sigma \notin U$. Setting k to the smallest integer with $\sigma^{p^k} \in U$ we conclude $k > 0$, $p^k \mid [G : U]$ and $p^{e+f-k} \mid \#U$. Thus, the claim is proven. \square

Remark 4.3. *In the case of a transitive group $G \subset S_n$ the above lemma can be applied to all cycles of p -power order having a fixed point. Then we can determine a divisor of $U := \text{Stab}_G(1)$. As this is a subgroup of index n . We can multiply the divisor of the order of $\#U$ by n and still get a divisor of $\#G$.*

4.2. Exclusion of block sizes. Is is well known, that the block size of a block system of a subgroup of S_n is a divisor of n . A priori, all divisors of n are possible block sizes. We use the following to exclude as many block sizes as possible. If we can exclude all of them, we have proved that the field is primitive. Thus, we can skip all the costly steps of the subfield search. As each permutation group with a block of size k is contained in a wreath product of the form $S_k \wr S_\ell$ we can use the following lemma.

Lemma 4.4. *Let $\sigma \in S_k \wr S_\ell$ be an element with cycle type (n_1, \dots, n_j) . Then, for each index $m = 1, \dots, j$ there is an integer e and an index set $I_m \subset \{1, \dots, j\}$ with $m \in I_m$ such that*

$$\forall i \in I_m : e \mid n_i \text{ and } \sum_{i \in I_m} n_i = e \cdot k.$$

Proof. Denote the blocks that contain the elements from the m -th orbit of σ by B_1, \dots, B_e . As σ acts transitively on the m -th orbit it has to act transitively on the blocks B_1, \dots, B_e as well.

Further, all elements in the blocks B_1, \dots, B_e are in orbits of size divisible by e . Let I_m be an indexing set for these orbits. We have $m \in I_m$ trivially. By counting the points in these blocks we get $\sum_{i \in I_m} n_i = e \cdot k$ and all summands are divisible by e . \square

Remark 4.5. *Lets inspect the particular case of a fixed point in the lemma above closer. This results in $e = 1$ and encodes that the fixed point is contained in a fixed block. Thus, the size of the block must be the sum of the length of some orbits of σ that includes the fixed point.*

5. INTERSECTION OF WREATH PRODUCTS

Remark 5.1. *As explained above, knowing a subfield of a number field is equivalent to knowing a block system of its Galois group. The latter means that the Galois group is contained in a certain wreath product.*

As soon as we know more than one subfield, we know that the Galois group is contained in the intersection of several wreath products. Thus, we need an efficient algorithm to compute the intersection.

Algorithm 5.2 (Intersection of wreath products). *Let $\mathcal{B}_1, \dots, \mathcal{B}_s$ be systems of blocks. Denote by n_1, \dots, n_s the number of blocks of the systems and by n the total number of points that is partitions into blocks.*

- (1) *Build a graph with a total number of $s + n_1 + \dots + n_s + n$ vertices.*
- (2) *Fix a bijection between the vertices and the systems of blocks, the blocks in all block systems and the n points.*
- (3) *Add an edge between a point-vertex and a block vertex if the point is contained in the block. Add an edge between a block vertex and a block systems vertex when ever the block is contained in the system of blocks.*
- (4) *Color the vertices representing the s block systems with s different colors.*
- (5) *Compute the group of color preserving automorphisms of the graph by using `nauty` [KP].*
- (6) *Compute the action of the automorphism group on the vertices representing the n points.*
- (7) *Return the image of the action as the intersection of the wreath products corresponding to the systems of blocks given.*

Remarks 5.3. (1) *To illustrate the advantage of the graph approach in contrast to the use of a general intersection algorithm, we construct permutation groups of degree 80 out of the 245 groups of order $17 \cdot 80$ by taking the coset-action with respect to a 17-Sylow subgroup.*

- (2) *We compute all principal systems of blocks of all these groups and build the corresponding wreath products. Using the standard intersection algorithm in MAGMA, we intersect all wreath products constructed. The total time for the intersections is 115 seconds. Using the graph-theoretic approach we can compute it in 6.5 seconds.*
- (3) *Doing the same in degree 165 with the 181 groups of order $8 \cdot 3 \cdot 5 \cdot 11$ by acting on the cosets of a 2-Sylow subgroup results in a total intersection time of 1251 seconds. The graph theoretic approach takes only 1.6 seconds.*
- (4) *A systematic test with all graphs that this approach constructs out of the database of transitive groups [Hul, CH] shows, that the back-track search of `nauty` never runs into a dead end.*
- (5) *In case we start Algorithm 5.2 with all principal systems of blocks of a permutation group and uses only one color in step 4, we get the normalizer of the intersection of the wreath products. When we try this, the run-time of `nauty` is much larger.*

- (6) *Because of all of this, we believe all the graphs that are constructed by Algorithm 5.2 to be special. We conjecture that computing the intersection of wreath products in this way is a polynomial time algorithm.*

6. GALOIS GENERATING SUBFIELDS

Definition 6.1. *The subfields $(K_i)_{i \in I}$ of L are called Galois-generating, if all block systems of the Galois group of L are block systems of the intersection of the wreath products corresponding to $(K_i)_{i \in I}$.*

Remark 6.2. *As we can express the intersection of subfields purely in terms of block systems, each intersection-generating set of subfields is Galois-generating as well. In many examples we observe that a minimal set of Galois-generating subfields is far smaller than a minimal set of intersection-generating subfields. We have to use this to compute the subfields efficiently.*

Example 6.3. *Let L be a field of degree $n = 2^m$ with elementary abelian Galois group C_2^m . Each pair of roots of f corresponds to a principal partition of block size 2. If we want to find all these principal subfields directly with the LLL-approach, we need $n - 1$ calls of the LLL-algorithm.*

However, when we know two of these principal subfields, the intersection K of these two fields results in a relative degree 4 extension L/K with two known subfields. Thus, the relative Galois group has to be the Kleinian four group. Thus, a third principal subfield is for free.

Calling the LLL once more gives us a new principal subfield with block size two. Inspecting each pair of the new and a known principal subfield in the same way as above gives us another 3 subfields for free.

This can be continued. We get all principal subfields out of m Galois-generating fields. Thus, m instead of $2^m - 1$ LLL-calls are sufficient.

- Remarks 6.4.**
- (1) *In the above example we used the structure of the Galois group to explain how new principal subfields are entailed by others. In general the combinatorics can be far more complicated. Thus, we take the detour via the block systems of the intersection of wreath products.*
 - (2) *Note that a system of blocks which is principal with respect to an intermediate group might not relate to a principal subfield. The point is, that a principal system of blocks for one group does not need to be principal with respect to all its subgroups.*
 - (3) *However, each group has at most $n - 1$ principal subfields and we derive at most $n - 1$ intermediate groups from the LLL-calls. Thus, in total at most $(n - 1)^2$ fields are inspected when we take all principal fields of all intermediate groups into account.*
 - (4) *We can terminate prematurely when we are sure to be finished.*
 - (5) *In practice many LLL-calls are performed to confirm that a known subfield is in fact principal with respect to a certain pair of roots. We can optimize the algorithm by detecting this out of the lattice formed by the subfields already known. Further, we can use the information that some subfields are proven to be principal.*

This leads to the following outline of the subfield algorithm:

Algorithm 6.5 (Outline of subfield algorithm). *To compute the subfields of a field L we do the following:*

- (1) *Factor the defining polynomial of L modulo various primes. Derive as much information as possible from the cycles found. In the best case this can prove that no subfields exist.*
- (2) *Select a prime p_s for a p_s -adic splitting field and a prime p for the LLL-computations.*
- (3) *Loop over the p -adic factors of f .*
- (4) *Test if the factor may result in a new principal subfield.*
- (5) *If a new principal subfield is possible, compute it and label it as proved to be principal.*
- (6) *If a new subfield is found make its system of blocks explicit in p_s -adic arithmetic and compute a Galois starting group out of all subfields known.*
- (7) *If the Galois starting group has more principal systems of blocks than subfields known compute the subfield polynomials and identify the first block in the p -adic arithmetic.*
- (8) *If we can derive from the Galois starting group that we are done, terminate prematurely.*

If we want all subfields explicitly, we can derive them from the subfields computed above as they contain at least all principal subfields. I.e., they are intersection generating.

7. DETAILED ALGORITHMS

To make the steps in the above outline more explicit, we give the field search algorithm and the sub-algorithms that are called.

Algorithm 7.1 (Field search).

Input: A polynomial f defining the number field $\mathbb{Q}[X]/f$.

Output: A list of intersection-generating subfields.

- (1) *Call Algorithm 7.2 for initialization. If this shows that no subfields exist terminate the algorithm.*
- (2) *Set up a p_s -adic splitting field of f and compute the p -adic factorization $f = f_1 \cdots f_m$ with $\deg(f_1) = 1$.*
- (3) *For each factor $f_j, j = 2, \dots, m$ do the following.*
- (4) *Call Algorithm 8.1 to check that there is space for an additional principal block system corresponding to f_j left. If not skip the factor.*
- (5) *Determine the principal subfield corresponding to f_j by calling Algorithm 7.6.*
- (6) *If a new subfield is found, compute the corresponding block system. Add the block system to the list of known block systems. Add a label to the field saying that it is proven to be principal.*
- (7) *If we have more than one known system of blocks, compute the group G as the intersection of the corresponding wreath-products by using Algorithm 5.2. If the Galois group is known to be even, intersect G with the alternating group.*
- (8) *If G has more principal block systems than known add these block systems and the corresponding subfields to the list of known block systems and subfields.*
- (9) *If the order of G coincides with the divisor of the group order known, terminate prematurely. The list of subfields known is intersection-generating.*

- (10) *In the case that the order of G is twice the lower bound of the group order, call Algorithm 7.3 and terminate the field search.*
- (11) *If all factors are inspected, return the list of subfields found as a list of intersection-generating subfields.*

Algorithm 7.2 (Prime inspection).

Input: A polynomial f defining the number field $\mathbb{Q}[X]/f$.

Output: ‘No subfields’ or

- Possible block sizes.
 - The LLL-prime p and the splitting field prime p_s .
 - A divisor of the order and the sign of $\text{Gal}(f)$.
- (1) *Enumerate the divisors of the degree of f as potential block sizes.*
 - (2) *Factor f modulo several primes. Store the cycle types found.*
 - (3) *Rule out all block sizes that are contradicted by Lemma 4.4 applied to the cycle types found. In the case that all block sizes are ruled out return ‘No subfields’.*
 - (4) *Count how many primes lead to at least one linear factor.*
 - (5) *Continue these steps, until at least one prime with a linear factor and a prime p_s with a reasonable splitting field degree are found.*
 - (6) *If all cycle types found correspond to even Frobenius permutations do the following: Test the discriminant of f to be a square. If so, note that the Galois group is even.*
 - (7) *Compute a divisor of the order of the Galois group from the cycle types found by using the methods described in Section 4.1.*
 - (8) *Out of the primes with a linear factor select the one with a minimal number of factors of degree less than the largest possible block size. Call this prime p .*
 - (9) *Let p_s be the inspected prime that results in the smallest p_s -adic splitting field.*
 - (10) *Return p and p_s as the LLL-prime and the splitting field prime we work with. Further, return the divisor of the group order and all block sizes that are not sieved out and the information whether the Galois group is even.*

The given field search algorithm is some type of compromise. It stops the LLL-steps if either all subfields are known or the Galois group is determined up to a factor of 2. At this point the idea is to look for the exact Galois group. Thus, all the index two subgroups of the known group have to be inspected with a quick test. This is done by the next algorithm.

Algorithm 7.3 (Final adjustment).

Input: $G \leq S_n$ with $[G : \text{Gal}(f)] \leq 2$, all subfield information known.

Output: Intersection generating subfields.

- (1) *Compute all index 2 subgroups of G .*
- (2) *Rule out the intransitive subgroups.*
- (3) *Rule out all subgroups with the same principal partitions as G .*
- (4) *For each subfield proven to be principal, rule out the subgroups that turn it into a non-principal subfield.*
- (5) *For each subgroup left, pick a new principal system of blocks and prove or disprove that it corresponds to an existing subfield by using Remark 7.5. (This can succeed for at most one subgroup.)*

- (6) *In the case that the last step was successful for one subgroup, compute all subfields corresponding to all new principal partitions of the new group and add them to the list of known subfields.*
- (7) *Return the list of subfields known as intersection-generating.*

Remark 7.4. *Example 7.8 and example 3 in Table 2 show the different possible behaviours of Algorithm 7.3. In the first case we descent to the right index 2 subgroup, in the second case we prove that the input was already the right starting group. Thus, in the first case we can reduce from 3 to 2 LLL-calls. In the second case we can skip 10 fruitless LLL-calls.*

Remark 7.5. *A method to prove that a conjectural block system is in fact a block system is described in [Kl, Algo. 44, 46]. The basic idea is to compute it by using Algorithm 3.9 and 3.10. The main difference is that we have to return fail in the case that either the coefficients of the subfield polynomial are bigger than predicted by the bound computed from the root bound or the coefficients of the embedding get to big. For the latter we can use the bound from [HKN, Lemma 18].*

Given two irreducible p -adic factors f_1, f_2 with $\deg(f_1) = 1$ of f , we have to compute the largest subfield L of K , such that the roots of f_1 and f_2 are in the same block with respect to the system of blocks corresponding to L .

The basic idea is to use the LLL-approach as described in [HKN] to construct it and use the same proof as above to confirm it. We give a detailed description how to merge the different approaches.

Algorithm 7.6 (Principal subfield).

Input: p -adic factors f_1, f_2 of f . p_s -adic roots r_1, \dots, r_n of f .

Output: The principal subfield to any root of f_2 .

- (1) *Choose an initial p -adic precision and lift the factors f_1, f_2 to this precision.*
- (2) *Build up the corresponding lattice as described in step 1 of algorithm Principal in [HKN].*
- (3) *Apply LLL-with-removals [HN, Algo. 2] to the lattice with the bound $n^2 \|f\|$. This results in a \mathbb{Q} -vector subspace $U \subset L$.*
- (4) *If U is of dimension 1, return \mathbb{Q} as principal subfield.*
- (5) *We denote by $h_1(\beta), \dots, h_k(\beta)$ the basis of U found.*
- (6) *Set the block identify precision σ to 1.*
- (7) *Compute the block values $V := \{(h_1(r_i), \dots, h_k(r_i)) : i = 1, \dots, n\}$ with p_s -adic precision σ .*
- (8) *If the set has less than k elements, double the precision σ and redo the last step.*
- (9) *If the set has more than k elements or one value occurs less than n/k time restart at step 2 with doubled precision.*
- (10) *Compute the potential block system*

$$\mathcal{B} := \{\{i \in \{1, \dots, n\} : (h_1(r_i), \dots, h_k(r_i)) = v\} : v \in V\}.$$
- (11) *Confirm the block system \mathcal{B} as described above. If this is successful, we get a defining polynomial of the subfield g and a root $h(\beta) \in L$ of it.*
- (12) *If the confirmation fails, restart at step 2 with doubled precision.*
- (13) *Check that the roots of f_1 and f_i are in the same block of the block system corresponding to the subfield found.*
- (14) *If the last step fails, restart at step 2 with doubled precision.*

Remark 7.7. (1) Note that the p -adic arithmetic in the LLL-computation and the p_s -adic arithmetic for the block identification and the subfield confirmation are independent. The interaction is only via U and the subfield polynomial g and its root $h(\beta)$. None of these data is of p -adic nature.

- (2) When we get the subspace U from the LLL-computation we can not assume that this is a subfield and we can not assume that it is the one we are targeting. We only know, that the subfield we are targeting is a vector subspace of U . Assuming U to be a subfield, we can identify the corresponding block system. In the case that the outcome is coarser than the expected block system (step 8) we increase the precision for the block identification.

If the outcome is finer than a block system (step 9), we have proved that U is not a subfield.

- (3) In the case that we successfully find a subfield, we are not yet done. If it does not relate to a block system which has the roots of f_1 and f_i in the first block the field found is larger than the one we target.

In the case that the last step is successful, we can conclude from the block analysis only that the subfield found is subfield of the targeted principal field. On the other hand U is a vector subspace that contains the targeted principal field. As the dimension of U and the subfield degree coincide we are done.

Example 7.8. Let K/k be a degree 60 Galois extension with group A_5 . As the extension is Galois all f_i are linear. The first call of the principal subfield algorithm gives us a degree 12 subfield corresponding to a block system of block size 5. The three remaining linear factors in the first block of this block system are skipped.

The second call of the principal subfield algorithm gives us a second degree 12 subfield that is a conjugate of the first one. Now, we compute the intersection of the corresponding wreath products. This is a group of order 240 isomorphic to an index 2 subgroup of $S_5 \times C_4$. It has 7 systems of blocks. As this group is not a subgroup of A_{60} , we can descent to the intersection U with A_{60} of order 120. U is isomorphic to $A_5 \times C_2$ and has 13 systems of blocks.

At this point, we determine all index two subgroups. As A_5 is simple, we find only one. It has 57 systems of blocks. After we confirmed one of the additional subfields by using Remark 7.5 we have proved that the extension is regular with Galois group A_5 . Thus, we are done.

8. THE LATTICE TEST

Let us assume the following situation. Let L, f be as above. Further, we have the factorization of f modulo p as $f_1 \cdots f_m$. We assume that f_1 is linear. We denote the root of f_1 by r_1 . In addition to this, we know already the subfields K_1, \dots, K_ℓ of K . We denote the block containing r_1 of the block systems corresponding to K_1, \dots, K_ℓ by $\Delta_1, \dots, \Delta_\ell$. As the Δ_i are Frobenius invariant, we can view each Δ_i as a subset of the above modulo p factorization of f . We want to check, if there is space for an additional principal subfield corresponding to the pair r_1, r_i . Here, r_i is a root of the local factor f_j . For this we use the algorithm of this section. We give the proof of correctness in Remark 8.2 and Lemma 8.3. Examples illustrating the efficiency are given in 8.4.

Algorithm 8.1 (Lattice test).**Input:**

- The p -adic factorization of $f = f_1 \cdots f_m$.
- The index $j \in \{2, \dots, m\}$ of the factor f_j to be tested.
- A list of possible block sizes.
- A list $\Delta_1, \dots, \Delta_\ell$ of known first blocks.
- Labels indicating which Δ_i are proven to be principal.

Output: ‘do factor’ or ‘skip factor’

- (1) Find the smallest first block Δ in $\Delta_1, \dots, \Delta_\ell$ that contains (the roots of) f_j .
- (2) In the case that no such first block is found set $\Delta := \{1, \dots, m\}$. (We are looking for a refinement of Δ .)
- (3) Set $n_0 := \sum_{i \in \Delta} \deg f_i$. This is the block size of Δ .
- (4) Compute the list N of all known refinements of Δ . These are all the Δ_i with $\Delta_i \subsetneq \Delta$.
- (5) In the case that $n_0 = 4$ and Δ is known to be a principal block and we have one refinement in N return ‘skip factor’.
- (6) In the case that $n_0 = 8$ and Δ is known to be principal and N contains a block of size 2 and a principal block of size 4, return ‘skip factor’;
- (7) From the list of potential block sizes extract the list S of all entries that are proper divisors of n_0 .
- (8) Delete from S all block sizes that can not be written as a sum of $1 + \deg f_j$ and the degrees of other factors in $\{f_i : i \in \Delta \setminus \{1, j\}\}$.
- (9) If a possible block size $d \in S$ satisfies $(d - \deg f_j) \cdot \frac{n_0}{k} < d$ with k the block size of a known refinement in N , delete d from S .
- (10) In the case that $n_0 = k \cdot q$ with a prime number $q > k$ and we know a refinement of Δ with block size q , delete q from S .
- (11) In the case that n_0 is the square of an odd prime number, do the following:
 - (a) In the case that there are two known refinements in N and Δ is known to be principal return ‘skip factor’.
 - (b) In the case that there are two known refinements in N and the factors $\{f_i : i \in \Delta \setminus \{1\}\}$ are not all of the same degree return ‘skip factor’.
- (12) In the case that $n_0 = p \cdot q$ with two prime numbers $q > p$ do the following:
 - (a) In the case that $q \not\equiv 1 \pmod p$ and N contains a refinement with block size p , delete p from S .
 - (b) In the case $p \cdot q \notin \{21, 55\}$ and a refinement with block size p is known do the following: If Δ is known to be principal or Δ contains a factor of degree > 1 , delete p from S .
 - (c) In the case that $p \cdot q \in \{21, 55\}$ and we know two refinements in N and Δ is known to be principal, return ‘skip factor’.
- (13) In the case that S is empty return ‘skip factor’.
- (14) return ‘do factor’.

Remark 8.2. The above test uses the following facts about block systems.

- (1) If a block system refines another block system the block size of the first one divides the block size of the second one.
- (2) Given two block systems, the non-empty intersections $B_i \cap B'_j$ are all of the same size.

- (3) In degree 4 the only permutation group with more than one (non-trivial) system of blocks is the Kleinian four group.
- (4) A degree 8 permutation group with principal blocks of size 2, 4, 8 has no other blocks.
- (5) In degree p^2 (p is an odd prime) a transitive group has either at most two systems of blocks or it is contained in $(C_p \times C_p) \rtimes C_{p-1}$. In the latter case it has $p+1$ systems of blocks and $\{1, \dots, p^2\}$ is not principal. Further, all cycles with a fixed point are of the form $(1, d, \dots, d)$ with a divisor d of $p-1$.
- (6) In degree $k \cdot q$ with $q > k$ (q a prime number) there is at most one system of blocks with block size q .
- (7) In degree $p \cdot q$ with $q > p$ (p, q both prime numbers) and $q \not\equiv 1 \pmod{p}$ there is at most one system of blocks with block size p .

Some of the above statements are easy. The p^2 -case is worked out in [DW]. As we do not know a good reference for the $(p \cdot q)$ -case we give the proof below.

Lemma 8.3. *Let p, q with $p < q$ be prime numbers and $G \subset S_{pq}$ be transitive with more than one block system of block size p . Then $q \equiv 1 \pmod{p}$.*

If G is solvable then G is the regular representation of $\mathbb{Z}/q\mathbb{Z} \rtimes \mathbb{Z}/p\mathbb{Z}$ with exactly q block systems of block size p . The trivial block system with only one block is not principal.

If G is not solvable we have $p \mid q-1 \mid p(p-1)$. Assuming the classification of finite simple groups we have one of the following two cases:

- (1) $p = 3, q = 7, \#G = 168$ and $G \cong \text{GL}_3(\mathbb{F}_2) \cong \text{PSL}_2(\mathbb{F}_7)$.
- (2) $p = 5, q = 11, \#G = 660$ and $G \cong \text{PSL}_2(\mathbb{F}_{11})$.

Both cases result in exactly two block systems of block size p and none of block size q .

Proof. Let $\mathcal{B}_1, \mathcal{B}_2$ be two block systems with block size p . Thus, we have $G \leq S_p \wr S_q$ and $q \mid \#G$. We denote the corresponding action on the blocks by φ_1 and φ_2 . The intransitive representation of G given by $\varphi_1 \oplus \varphi_2$ is faithful. Thus, G is a sub-direct product in $\varphi_1(G) \times \varphi_2(G)$.

As $\#G, \#\varphi_1(G)$ and $\#\varphi_2(G)$ are exactly once divisible by q and each nontrivial normal subgroup of a transitive permutation group of prime degree q contains all the q -Sylow-subgroups [Hup, Kap. II Satz 1.5] we can conclude that φ_1 and φ_2 are faithful permutation representations of G .

In the case that $G \cong \varphi_1(G)$ is solvable we can use the result of Galois $\varphi_1(G) \cong (\mathbb{Z}/q\mathbb{Z}) \rtimes (\mathbb{Z}/r\mathbb{Z})$ [Hup, Kap. 2, Satz 3.6]. Here, r is a divisor of $q-1$. As G acts transitive on $p \cdot q$ points we conclude $p \mid r \mid q-1$. In the case that $r = p$ the group G has exactly q Sylow p -subgroups and one Sylow q -subgroup. Thus, we have q block systems with block size p and one with block size q . If $p < r$ the point stabilizer in G is a cyclic group of order $\frac{r}{p}$. The unique index q subgroup that contains it is its normalizer. Thus, G has only one block system of block size p .

Now we assume that $G \cong \varphi_1(G)$ is not solvable. By a theorem of Burnside [Hup, Kap. 5, Satz 21.3] $\varphi_1(G)$ is 2-transitive. Let $U \subset G$ be the stabilizer of one block of \mathcal{B}_1 . The orbits of $\varphi_1(U)$ are of length 1 and $q-1$. Thus, U has one orbit of length p and the lengths of the other orbits of U are multiples of $q-1$.

The orbit of length p consists of one block of \mathcal{B}_1 . It hits p blocks of \mathcal{B}_2 . These blocks form a U -invariant set of order p^2 . Thus, p^2 is the sum of p and some

multiples of $q - 1$. This shows $(q - 1) \mid (p^2 - p) = p(p - 1)$. As q is bigger than p we can conclude $p \mid q - 1 \mid p(p - 1)$.

As $G \cong \varphi_1(G)$ is of prime degree q and not solvable the classification of finite simple groups implies that $\varphi_1(G)$ is one of: [DM, Sec. 3.5]

- 1) The symmetric or the alternating group of degree q .
- 2) $\text{PSL}_2(\mathbb{F}_{11})$ acting on $q = 11$ points.
- 3) The Mathieu group M_q with $q = 11$ or $q = 23$.
- 4) A projective linear group G with $\text{PSL}_d(\mathbb{F}_\ell) \subset G \subset \text{PGL}_d(\mathbb{F}_\ell)$ of degree $q = \frac{\ell^d - 1}{\ell - 1}$. Here, d is a prime.

Further, we have proved that the one-point stabilizer of $\varphi_1(G)$ has an index p subgroup with $p \mid q - 1 \mid p(p - 1)$. First, this implies $q \in \{3, 7, 11, 23, 43, 47, 53, \dots\}$. As S_3 is solvable, q is at least 7. Now we check, which of the above options is compatible with this.

- 1) The one-point stabilizer of the symmetric and the alternating groups of degree $q \geq 7$ is a the symmetric or the alternating group of degree $q - 1$. The subgroups of smallest index have index 2 or $q - 1$. They are given by the intersection with the alternating group or as the stabilizer of a second point. Further, if $q - 1 = 6$ we get another index 6 subgroup. This excludes the option of a suitable index p subgroup.
- 2) The one-point stabilizer of $\varphi_1(G) = \text{PSL}_2(\mathbb{F}_{11})$ has exactly one maximal subgroup of prime degree. It results in one of the exceptions listed in the claim.
- 3) The one-point stabilizers of the Mathieu groups have only one subgroup of prime index. This is an index 2 subgroup in the one-point stabilizer of M_{11} . This excludes the option.
- 4) Here we have $p \mid \frac{\ell^d - 1}{\ell - 1} - 1 \mid p(p - 1)$. In the special case of $d = 2$ we get $p \mid \ell \mid p(p - 1)$. As p is a prime number and ℓ a prime power we conclude $p = \ell$ and $q = p + 1$. This forces $p = 2, q = 3$. This contradicts $q \geq 7$.

Otherwise $d = 2k + 1$ is an odd prime number. This results in

$$\begin{aligned} p \mid q - 1 &= \ell(\ell^{2k-1} + \ell^{2k-2} + \dots + \ell + 1) \\ &= \ell(\ell^k + 1)(\ell^{k-1} + \ell^{k-2} + \dots + \ell + 1) \mid (p - 1)p. \end{aligned}$$

As $q - 1$ is larger than $\ell^{2k} + \ell^{2k-1}$ we get the estimate

$$\ell^{2k} + \ell^{2k-1} < q - 1 \leq p(p - 1) < p^2.$$

This implies $p \geq \ell^k + 1$. Here, equality is only possible in the case of $k = 1$. As the largest factor of $q - 1$ is $\ell^k + 1$, we can conclude $p = \ell^k + 1, k = 1$, and $d = 3$.

Using once more that ℓ is a prime power we get that p is a Fermat prime and ℓ has the shape 2^{2^e} . As $q = \ell^2 + \ell + 1$ we can conclude $3 \mid q$ as long as 2^e is even. Thus, the last remaining possibility is $e = 0, \ell = 2, p = 3, q = 7$, and $d = 3$. It results in $\varphi_1(G) = \text{GL}_3(\mathbb{F}_2)$ as a potential group. This is the other exception in the claim. \square

Examples 8.4. When applying the lattice test to a root r_i of a local factor f_j of the field extension K/\mathbb{Q} the algorithm will first determine the largest known subfield k such that r_1 and r_i are in the same block. Then our criteria will be applied to

the relative extension K/k . To keep the examples clear, we will describe only the relative situation:

- 1) Assume that K/k is a cyclic extension of degree 15. As the Galois group is a regular permutation group, we get 15 linear factors f_j . We need two successful calls to the principal subfield algorithm to get the subfields of degree 3 and 5. After this, the $p \cdot q$ -case of the lattice test applies and stops the search for further principal subfields. Note, that the intersection of the wreath products will result in the group $S_3 \times S_5$. Using the discriminant we can descend to an index 2 subgroup.
- 2) Let K/k be a cyclic extension of degree 10. As above, all f_j are linear. Here we have to find the degree 2 and the degree 5 subfield with one successful call of the principal subfield algorithm. Further, we have to confirm that k is principal. Otherwise the Galois group could be the regular representation of a dihedral group with 4 other subfields. Thus, in the lucky case 3 LLL-calls suffice.
- 3) Assume that K/k is a Galois extension of degree 10 with a dihedral Galois group. Here we have to determine two block systems with block size 2. Then the intersection of the wreath products descends to a group of order 10 and gives us 4 other subfields. Knowing all these subfields the lattice test will stop the search. Aside from this, the lower bound of the group order could be used.
- 4) Let K/k be a degree 21 extension with Galois group $21T4 \cong C_7 \rtimes C_6$. Here, we can find a prime with three linear and nine quadratic local factors. A first call of the principal subfield algorithm with a linear local factor will result in a degree 7 subfield. Now, the third linear factor can be skipped as it is part of a known block of size 3.

A successful call of the principal subfield algorithm with a quadratic factor gives us a block system of block size 7 corresponding to a cubic subfield. At this point the $p \cdot q$ test applies. It concludes that the only possibility of even more subfields corresponds to the regular Galois group $C_7 \rtimes C_3$. As we started with an irregular cycle, this option is excluded. Thus, we have all subfields.

- 5) Assume K/k to be a degree 8 extension with Galois group $8T17$ of order 32. We can choose a prime such that four of the f_i are linear and the fifth is quartic.

As no block system with block size bigger than 4 is possible, we can skip the quartic factor and label k as principal. Applying the principal subfield algorithm twice to linear factors gives us a quadratic and a quartic principal subfield. Now, we can skip the last linear factor. This follows from the statement on degree 8 permutation groups as we have principal subfields of degree 1, 2, 4.

9. APPLICATIONS TO GALOIS GROUP COMPUTATION

The computation of the Galois group of the splitting field of a polynomial f of degree n with rational coefficient is usually done by a method introduced by Stauduhar [St]. It constructs a descending chain of subgroups starting at a sufficiently large group (e.g. S_n) down to the Galois group.

To perform well in practice, we have to start with a group as small as possible. Subfields are used at this point, as each of them relates to a wreath product that contains the Galois group. Initially, the Galois group of each subfield was determined, to get an even smaller wreath product. Finally the intersection of all wreath products was determined [Ge, Algorithmus 5.3].

Doing it this way results in two disadvantages. First, non-maximal subfields are treated multiple times. Once for each field containing it.

Further, the intersection of several wreath products can already result in a very small group. Thus, the strategy should be rearranged. We present it as an algorithm. We remark that we perform step (2) using the graph theoretic methods from Section 5. We have the hope that this step is in polynomial time, but we can not prove this.

Algorithm 9.1 (Galois starting group).

Input: A field L .

Output: An overgroup of the Galois group of L .

- (1) *Determine Galois-generating subfields of L .*
- (2) *Compute the intersection of the wreath products corresponding to these fields. Call this group G_0 .*
- (3) *Determine the projection of G_0 to a starting group of the Galois group of $\mathbb{Q}(\{\sqrt{\text{Disc}(K_i)} : i \in I\})/\mathbb{Q}$. Here, $\{K_i : i \in I\}$ is the set of all subfields of L .*
- (4) *Determine the Galois group of the multi-quadratic extension above and replace G_0 by the pre-image of this group.*
- (5) *Find the subfield K of smallest degree such that the projection of G_0 to a Galois starting group of K does not result in a group of order equal to the lower bound of the Galois group of K computed by cycle type inspection.*
- (6) *Determine the projection π_K that maps G_0 to its action on the block system corresponding to K .*
- (7) *Use $\pi_K(G_0)$ as a starting group to determine the group G_K of K by the Stauduhar method.*
- (8) *Replace G_0 by $\pi_K^{-1}(G_K)$.*
- (9) *Redo the last three steps for all other subfields K with degrees in ascending order.*
- (10) *If L/\mathbb{Q} has a unique maximal subfield K and $[L : K] > 2$ compute the field $K_\Delta := K[\text{Disc}(L/K)]$ and the homomorphism π that maps G_0 to a starting group for K_Δ .*
- (11) *In the case that K_Δ was constructed, determine the Galois group G_Δ by the Stauduhar method and replace G_0 by $\pi^{-1}(G_\Delta)$.*
- (12) *Return G_0 as a starting group for the Galois group computation of L .*

Remark 9.2. *In step 10 we use another auxiliary field that is not part of the subfield lattice. This is only done in the case that the relative degree is bigger than 2 and the maximal subfield is unique. The reason for these restrictions is that otherwise this subfield would coincide with L or the intersection of the wreath products together with the subfield inspection results already in a small starting group. Thus, in the latter case there is no reason to expect a large index descent in step 11.*

A systematic test with the database of transitive groups brings the degree 42 group with number 5798 to light. It relates to a degree 42 field with maximal subfields of

degree 14 and 21. The intersection of the two maximal subfields is of degree 7. The subfield groups are $S_2 \wr S_7$ and $S_3 \wr S_7$. Thus, in this case the computed starting group is $D_{2 \cdot 6} \wr S_7$. Here, step 10 would construct a degree 28 field with starting group $(S_2 \times S_2) \wr S_7$. Applying step 11 would result in an descent to a maximal subgroup of index 64.

Example 9.3. Lets illustrate this strategy by applying it to $f_{18} = x^{18} + 9x^9 + 27$. The subfield algorithm determines the subfields given by $f_2 := x^2 + 9x + 27$, $f_3 := x^3 - 12x^2 + 39x - 37$, and $f_6 := x^6 + 9x^3 + 27$. The sextic field is the composition of the smaller ones. The subfields correspond to the systems of blocks

$$\begin{aligned} & \{\{1, 2, 3, 4, 5, 6, 7, 8, 9\}, \{10, 11, 12, 13, 14, 15, 16, 17, 18\}\}, \\ & \{\{1, 2, 3, 13, 14, 15\}, \{4, 5, 6, 16, 17, 18\}, \{7, 8, 9, 10, 11, 12\}\}, \\ & \{\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8, 9\}, \{10, 11, 12\}, \{13, 14, 15\}, \{16, 17, 18\}\}. \end{aligned}$$

The intersection of the corresponding wreath products results in the group 18T903 of order 559872. This projects to a starting group for the sextic subfield of order 12. As the discriminant of f_3 is a square we can pass to an index 2 subgroup. This determines the groups of all subfields.

As f_2 and f_{18} have the same discriminant, we can do another descent to an index 2 subgroup. This gives us the next smaller starting group 18T806.

At this point the algorithm determines that the Galois group of $\mathbb{Q}[X]/f_{18}$ viewed as a degree 3 extension of $\mathbb{Q}[X]/f_6$ is even as the resolvent $x^{12} + 531441x^6 + 282429536481$ is reducible. Thus, we return the group 18T453 of order 4374 as starting group for the Galois group computation.

In the case that we want to finish the computation of the Galois group of f_{18} , we descent via 18T323, 18T160 and 18T82 to 18T16. All descents are of index 3.

10. ALGORITHM SELECTION

At a first glance it seems that the non-polynomial time algorithm [Kl] is now outdated. However, this is not the case. The point is, that there are many cases, where the old Algorithm had to do only very few steps, each of them being a lot simpler than an LLL-reduction.

This can be explained by interpreting the old algorithm as part of a Galois group computation with the Stauduhar method. This method determines the Galois group by constructing a descending chain of overgroups of the Galois group till the Galois group is reached.

More precisely, determining a subfield of degree ℓ is equivalent to determining a wreath product $S_k \wr S_\ell \subset S_n$ that contains the Galois group. The number of subgroups conjugate to $S_k \wr S_\ell$ is $\binom{n-1}{k-1} \binom{n-k-1}{k-1} \dots \binom{2k-1}{k-1} \binom{k-1}{k-1}$.

Testing all these conjugates is too costly. The idea used in [Kl] is, that most of these conjugates do not contain the known Frobenius element of the Galois group. In other words, the Frobenius does not respect most decompositions of roots into blocks.

In the more general setting of Galois group computations the conjugate subgroups of a subgroup that contain a prescribed Frobenius element are determined by using *short cosets* which are formed by the coset representatives σ of G/U such that $\text{Frob} \in U^\sigma$ [Ge]. There are several ways to determine short cosets efficiently [E1].

The number of short cosets can be derived from the cycle type of the Frobenius permutation. In table 1 we give a few examples for a regular cycle.

Cycle type of Frobenius	Possible cycle types on blocks	Number of short cosets
(n)	(ℓ)	1
$(\frac{n}{2}, \frac{n}{2})$	$(\ell), (\frac{\ell}{2}, \frac{\ell}{2})$	$\ell, 1$
$(\frac{n}{3}, \frac{n}{3}, \frac{n}{3})$	$(\ell), (\frac{\ell}{3}, \frac{2\ell}{3}), (\frac{\ell}{3}, \frac{\ell}{3}, \frac{\ell}{3})$	$\ell^2, 2\ell, 1$
$(\frac{n}{4}, \frac{n}{4}, \frac{n}{4}, \frac{n}{4})$	$(\ell), (\frac{\ell}{2}, \frac{\ell}{2}), (\frac{\ell}{4}, \frac{3\ell}{4}),$ $(\frac{\ell}{4}, \frac{\ell}{4}, \frac{\ell}{2}), (\frac{\ell}{4}, \frac{\ell}{4}, \frac{\ell}{4}, \frac{\ell}{4})$	$\ell^3, \frac{3}{4}\ell^2, \frac{9}{4}\ell^2,$ $3\ell, 1$
$(1, \dots, 1)$	$(1, \dots, 1)$	$\frac{n!}{n(n-k)(n-2k)\dots(2k)k}$

TABLE 1 – Number of short cosets for $S_k \wr S_\ell \subset S_{k\ell}$

The table shows that the algorithm described in [Kl] should be used in the case that a Frobenius element with a small number of orbits is found.

11. PRACTICAL PERFORMANCE TEST

All test are carried out by using MAGMA 2.22 on one core of an intel i5-4690 processor running at 3.5 GHz. For testing the method we use the test polynomials listed in <http://www.math.fsu.edu/~hoeij/subfields/>. This list covers a range of cases. They all have in common that the number of p -adic factors is not too small. The performance of the algorithms are listed in Table 2.

Nr	Degree	Divisor of #Gal	Galois group	# G_0	LLL calls	LLL time	total time	time for [Kl]
1	36	108	$C_3 \times S_3 \times S_3$	108	10	1.73	2.30	0.43
2	75	300	$C_5^2 \rtimes C_{12}$	600	8	84.55	86.88	26.39
3	48	192	$C_4^2 \rtimes D_{2.6}$	384	7	4.85	6.09	6.16
4	56	336	$\text{Aut}(\text{GL}_3(\mathbb{F}_2))$	336	3	5.34	6.03	1415.99
5	50	100	$C_5^2 \rtimes C_4$	200	3	3.21	3.82	99.19
6	60	120	S_5	120	5	14.12	16.81	98.07
7	64	576	$C_2^4 \rtimes (S_3 \times S_3)$	2304	19	84.19	86.44	105.47
8	72	288	$C_6^2 \rtimes D_{2.4}$	288	6	16.47	20.84	3934.64
9	60	60	A_5	120	2	9.5	16.36	2442.16
10	81	162	$C_3 \times (C_3^3 \rtimes C_2)$	324	5	67.90	79.09	7405.01
11	81	162	$C_3^2 \wr S_2$	324	8	124.65	133.71	13945.61
12	32	32	C_2^5	32	5	0.41	6.62	> 50h
13	64	256	$D_{2.4} \times (C_4 \rtimes D_{2.4})$	512	9	26.38	33.78	> 50h
14	96	11520	$S_6 \times C_2^4$	11520	6	111.46	142.52	> 50h
15	96	96	$C_2^4 \rtimes C_6$	96	3	61.07	107.31	42395.95
16	75	300	$C_5^2 \rtimes (C_3 \rtimes C_4)$	600	3	72.29	77.65	93621.32
17	80	160	$C_2^4 \times D_{2.5}$	320	6	90.50	113.14	> 50h
18	90	360	A_6	720	8	284.23	290.03	> 50h
19	100	100	Equal to Nr. 5	100	3	188.97	224.06	> 50h
20	64	64	C_2^6	64	6	27.01	339.92	> 50h

TABLE 2 – Performance of the subfield algorithms (all times in seconds)

We can read from the table that except for the first two examples, the new method is faster. Further, the lower bound of the order of the Galois group is exact in all examples. The row with $\#G_0$ lists the order of the Galois starting group in the moment we stop calling the principal subfield algorithm.

Compared to the table given in [HKN] our hardware seems to be about a factor 2.5 faster, as the run times of the unchanged implementation of [Kl] are about that shorter. For the LLL-based approach, our computation is a factor of $4.8 - 190$ faster. Taking the newer hardware into account, the implementation is a factor $1.9 - 75$ better. To a small part this improvement can be explained by the faster subfield proof and a potentially better LLL implementation. The main reason for the speed up is the reduced number of LLL-calls. In the degree 64 example Nr. 7 we have the smallest improvement. The number of p -adic factors is 24. We have to call the LLL for 19 of them. In the degree 100 example, we have 100 p -adic factors but the LLL is called only 3 times. This explains why we have the biggest improvement in this example.

To go even further, we would need a better strategy to select the p -adic factor treated next and a good heuristic to stop the principal subfield search even earlier and finish with another (e.g. the Stauduhar) method.

REFERENCES

- [BCP] W. Bosma, J. Cannon, and C. Playoust: *The Magma algebra system. I. The user language*. J. Symbolic Comput. **24** (1997), 235 – 265.
- [BFSS] A. Bostan, P. Flajolet, B. Salvy, E. Schost: *Fast computation of special resultants*. J. Symbolic Comput. **41**, no. 1 (2006), 1 – 29.
- [CH] J. Cannon, D.F. Holt: *The transitive permutation groups of degree 32*. Experiment. Math. **17** (2008), no. 3, 307 – 314.
- [DM] J. Dixon, B. Mortimer: *Permutation groups*. Springer-Verlag, New York 1996.
- [DW] E. Dobson, D. Witte: *Transitive permutation groups of prime-squared degree*. J. Algebraic Combin. **16** (2002), no. 1, 43 – 69.
- [E1] A.S. Elsenhans: *A note on short cosets*. Exp. Math. **23** (2014) no. 4, 411 – 413.
- [E2] A.S. Elsenhans: *Improved methods for the construction of relative invariants for permutation groups*. J. Symbolic Comput. **79** (2017), part 2, 211 – 231
- [Ge] K. Geißler: *Berechnung von Galoisgruppen Über Zahl- und Funktionenkörpern*. Dissertation, Berlin, 2003.
- [HKN] M. van Hoeij, J. Klüners, A. Novocin: *Generating Subfields*. J. Symbolic Comput. **52**, 2013, 17 – 34.
- [HN] M. van Hoeij, A. Novocin: *Gradual Sub-lattice Reduction and a New Complexity for Factoring Polynomials*. Algorithmica **63** (2012) no. 3, 616 – 633.
- [HEB] D.F. Holt, B. Eick, E.A. O'Brien: *Handbook of computational group theory*. Chapman & Hall/CRC, Boca Raton, FL, 2005.
- [Hul] A. Hulpke: *Constructing transitive permutation groups*. J. Symbolic Comput. **39** (2005), no. 1, 1 – 30.
- [Hup] B. Huppert. *Endliche Gruppen I*. Berlin-New York 1967.
- [Kl] J. Klüners: *On Computing Subfields - A Detailed Description of the Algorithm*. Journal de Theorie des Nombres de Bordeaux **10** (1998), 243 – 271.
- [KP] B. McKay, A. Piperno: *Practical graph isomorphism, II*. J. Symbolic Comput. **60** (2014), 94 – 112.
- [LLL] A. K. Lenstra, H. W. Lenstra, L. Lovász: *Factoring polynomials with rational coefficients*. Math. Ann. **261** (1982), no. 4, 515 – 534.
- [St] R. Stauduhar: *The determination of Galois groups*. Math. Comp. **27** (1973), 981 – 996.